

ORACLE

Driving Success through Innovation and Trust

Jungok Nah

Vice President, Head of Cloud Engineering, Oracle Korea

Feb 11, 2025





70%

of workloads are still
on-premises

72%

of organizations
adopted GenAI in at
least one business
function

50%

of multinationals will
have a digital
sovereignty strategy
by 2028

**Adopt
Cloud**

**Use
AI**

**Control
Data**

**Distributed
cloud**

**AI
Infrastructure**

**Distributed
AI**

OCI Distributed Cloud

The full cloud where and how you need it

OCI Distributed Cloud



Public cloud

Access cloud services from global regions



Multicloud

We work well with other clouds



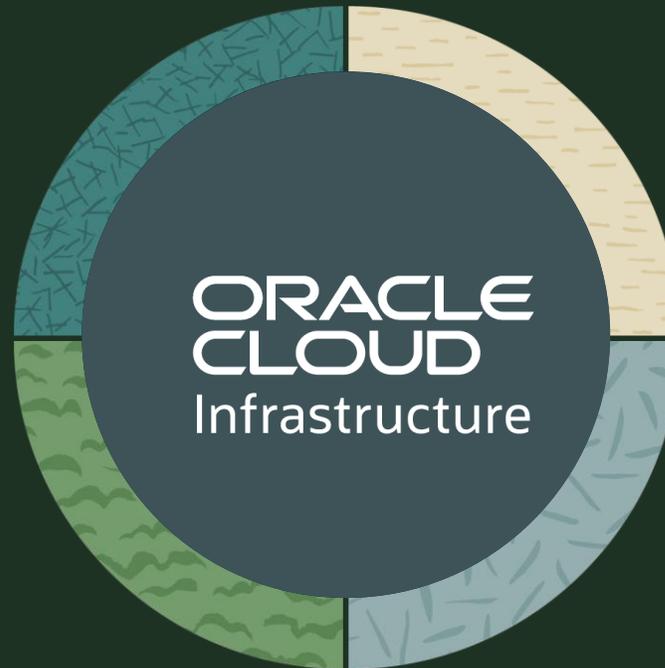
Hybrid/Edge

OCI brings cloud services to your choice of locations



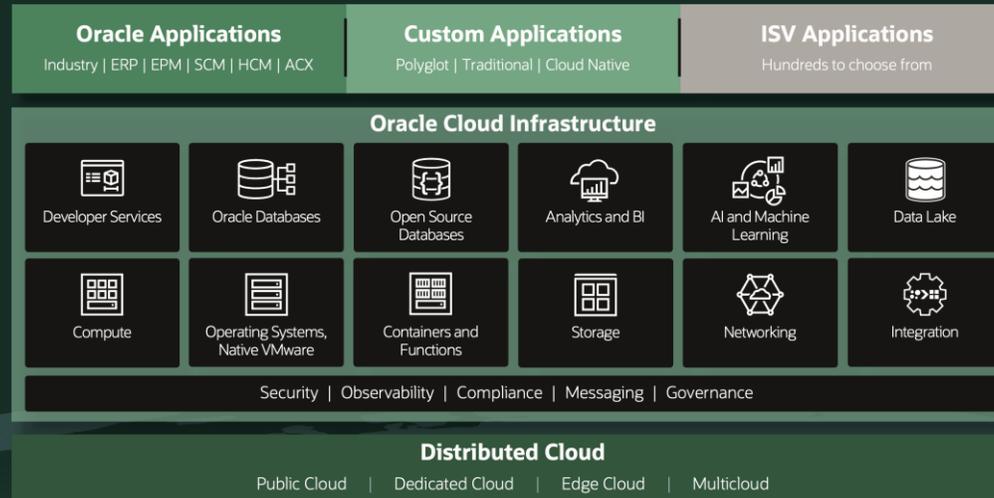
Dedicated cloud

We build a cloud just for you



Everything we offer is available everywhere

For you to focus on innovation, not cloud complexity



Public Cloud



Gov Cloud



Isolated
Regions



Cloud@
Customer



Dedicated
Regions



Multicloud



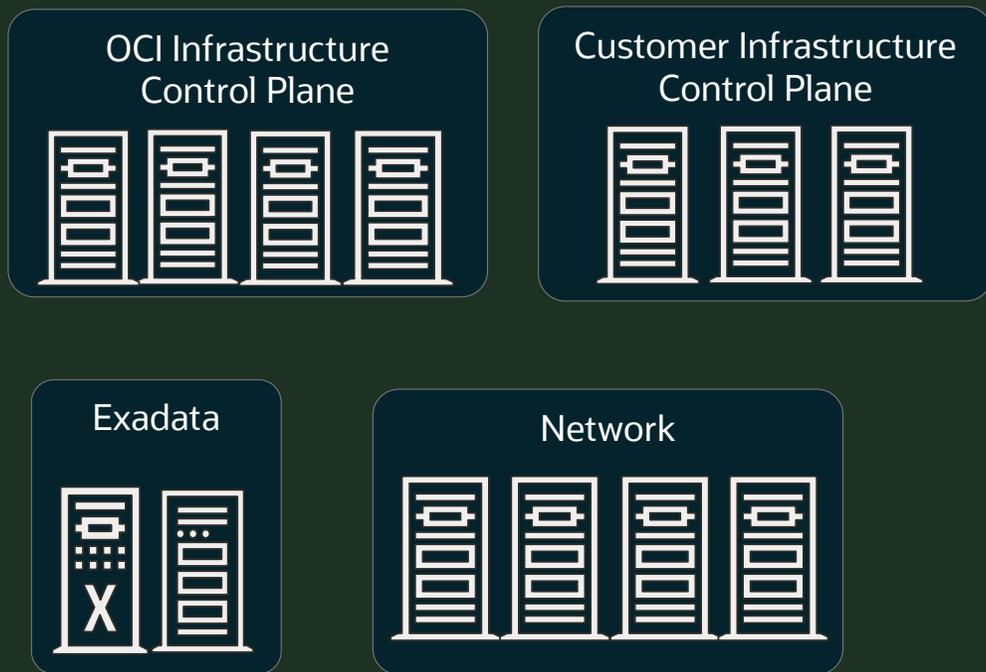
Alloy



EU Sovereign



The Cloud: in as little as **3** racks



Dedicated Region current footprint : 13 racks, 126 hosts

Dedicated Region²⁵

Full cloud	Flexible sizing	Easier than ever
------------	-----------------	------------------



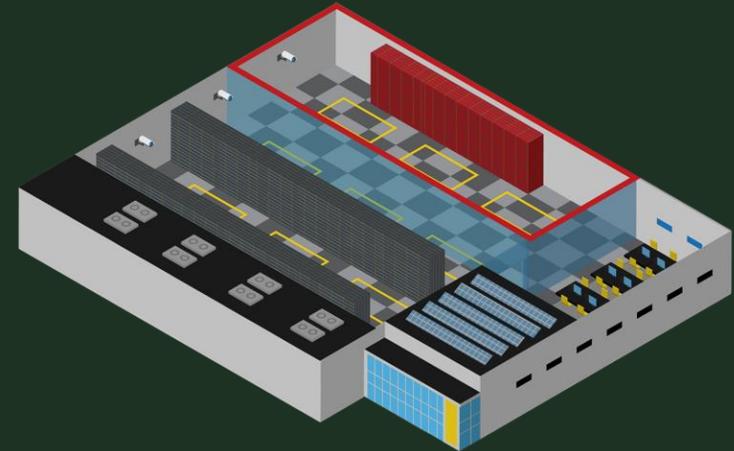
3 Hyper Converged Racks
36 Hyper Converged Servers



Open Multi Cloud Era: OCI Distributed Cloud

ORACLE

The world's best database is now available at all hyperscalers



Cloud Regions Everywhere

- *Public cloud*
- *Government cloud*
- *EU Sovereign cloud*
- *Isolated Region*
- *Multicloud*
- *Dedicated cloud on-prem*

Oracle AI Infrastructure

The most powerful AI infrastructure

Core tenets for AI Infrastructure



Scale



Performance



Speed

Most powerful AI infrastructure



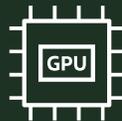
ONLY
hyperscaler with
bare metal GPUs
for AI training



Purpose-built
RoCEv2
networking at
104 Pb/s



High performance
file storage with
100x
performance
of standard
file storage



131,072 GPUs
in one OCI Supercluster



4 of the largest
LLMs run on OCI

AI infrastructure for all your needs

		Now GA				Now Orderable	Now orderable	Now orderable	Now orderable
GPU.A10.1 GPU.A10.2 GPU.A10.4 VM, bare metal 1 x A10 24 GB 2 x A10 24 GB 4 x A10 24 GB	Coming soon GPU.A100.1 GPU.H100.1 VM 1 x A100 80 GB 1 x H100 80 GB	GPU.L40S.4 Bare metal 4 x L40S 48 GB 192 GB GPU memory 7.7 TB Local NVMe 200 Gbps Front-end network	GPU.4.8 Bare metal 8 x A100 40 GB 320 GB GPU memory 27.2 TB Local NVMe 100 Gbps Front-end network	GPU.A100-v2.8 Bare metal 8 x A100 80 GB 640 GB GPU memory 27.2 TB Local NVMe 100 Gbps Front-end network	GPU.H100.8 Bare metal 8 x H100 80 GB 640 GB GPU memory 61.4 TB Local NVMe 100 Gbps Front-end network	Now Orderable GPU.MI300X.8 Bare metal 8 x MI300X 192 GB 1.5 TB GPU Memory 30.7 TB Local NVMe 100 Gbps Front-end Network	Now orderable GPU.H200.8 Bare metal 8 x H200 141 GB 1.1 TB GPU memory 30.7 TB Local NVMe 200 Gbps Front-end network	Now orderable GPU.B200.8 Bare metal 8 x B200 192 GB 1.5 TB GPU memory 30.7 TB Local NVMe 400 Gbps Front-end network	Now orderable GPU.GB200.4 NVL72 (18 nodes/rack) Bare metal 36 Grace CPUs 72 Blackwell GPUs 13.5 TB GPU memory 30.7 TB GPU+CPU Memory 553 TB Local NVMe 7,200 Gbps Front-end network

Object storage and/or File Storage	Now GA File Storage HPMT—20/40/80 Gb/sec per MT	Coming soon Managed Lustre—8 Gb/sec per TB
------------------------------------	--	---

SMALL-SCALE

MEDIUM-SCALE

LARGE-SCALE AI TRAINING, INFERENCE, FINE-TUNING

OCI Superclusters	3,840 GPUs	32,768 GPUs	32,768 GPUs	16,384 GPUs	16,384 GPUs	65,536 GPUs	131,072 GPUs	100,000+ GPUs
	800 Gbps RDMA	1,600 Gbps RDMA	3,200 Gbps RDMA	28,800 Gbps RDMA				



Distributed AI

Your data is distributed, Bring AI to your data

“Everything Everywhere”

All 100+ OCI Services and up to 80+ compliance certifications

Any organization

Government regulated use

Oracle-operated data center

Your data center

Cloud services for your organization

Cloud services for resale to your customers

Any workload

Workloads restricted to a geopolitical area

Workloads restricted to government usage

Workloads requiring air-gapped high security

Global operations
Authorized personnel

Geographical ops
Authorized personnel

Restricted ops
Cleared personnel

Disconnected ops
Cleared personnel



Commercial regions

38 Global regions, and growing



Sovereign regions

2 EU regions



Government regions

8 US, UK, and Australia regions



Oracle National Security regions



Dedicated Region



Oracle Alloy

Provide cloud services with your logo to your end customers



Isolated Region

Dedicated Region and Oracle Alloy are not restricted to any particular region. They can be deployed for multiple use cases.



OCI Distributed Cloud

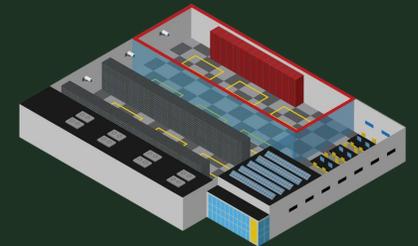
The full cloud where and how you need it



Same Great Price
Everywhere



Built-in Security
Everywhere



Cloud Regions
Everywhere



The World's AI is being built on OCI

Join us

ORACLE